arXiv:1607.02559v1 [cs.LG] 9 Jul 2016

Uncovering Locally Discriminative Structure for Feature Analysis*

Sen Wang[1], Feiping Nie[2], Xiaojun Chang[3]
Xue Li[1], Quan Z. Sheng[5], and Lina Yao[4]

[1] School of Information Technology and Electrical Engineering,
The University of Queensland, Australia.
[2] Centre for OPTIMAL, Northwestern Polytechnical University, Shaanxi, China.
[3] Centre for Quantum Computation & Intelligent Systems,
University of Technology Sydney, Australia.
[4] School of Computer Science, The University of Adelaide, Australia.
[5] School of Computer Science and Engineering, The University of New South Wales, Australia.

Abstract. Manifold structure learning is often used to exploit geometric information among data in semi-supervised feature learning algorithms. In this paper, we find that local discriminative information is also of importance for semi-supervised feature learning. We propose a method that utilizes both the manifold structure of data and local discriminant information. Specifically, we define a local clique for each data point. The k-Nearest Neighbors (kNN) is used to determine the structural information within each clique. We then employ a variant of Fisher criterion model to each clique for local discriminant evaluation and sum all cliques as global integration into the framework. In this way, local discriminant information is embedded. Labels are also utilized to minimize distances between data from the same class. In addition, we use the kernel method to extend our proposed model and facilitate feature learning in a high-dimensional space after feature mapping. Experimental results show that our method is superior to all other compared methods over a number of datasets.

1 Introduction

The performance of machine learning tasks, e.g. classification or clustering, is mainly affected by the input features that are extracted from raw data. Learning distinctive features or effective data representations can without doubt benefit the consequent learning tasks. Over the past decade, feature analysis has attracted much research attention in different fields, such as machine learning [32,43,33,40], multimedia analysis [33], biomedical applications [45,42,44], etc. In literature, a number of unsupervised and supervised methods have been developed to learn new features. Clustering algorithms have been widely used as an unsupervised feature learning procedure to obtain statistical data representation. A typical example in multimedia analysis is k-means, which learns a dictionary from a number of image or video training samples without class information [34,21]. Even though increasing the size of the dictionary can squeeze out a

* Corresponding Author: Feiping Nie. Email: feipingnie@gmail.com

bit of extra performance, it is still difficult to identify the best choice of the number of centers.

As one of the representative supervised algorithms, Linear Discriminant Analysis (LDA) [8] finds the best data projection maximizing distances between different class centers while making data samples from the same class closer to each other. This is achieved by maximizing the ratio of the between-class covariance to the within-class covariance. When there are sufficient labeled training data, LDA-based features effectively support machine learning algorithms in a variety of applications. For example, LDA has been used to strengthen the class information of face images in many face recognition systems [11,31,2,38,29,5,24,23]. Also, LDA has been revisited in [1] and been evaluated in three pipelines over a few face image datasets for the purpose of gender recognition. In [17], authors evaluate three LDA-based variants to obtain a discriminant movement representation for multi-view action recognition. Unfortunately, when the number of training samples is small, LDA suffers from the *small sample size* (SSS) problem. This is because the small number of training samples will make the within-class scatter matrix singular, which would result in computational difficulty. Meanwhile, learning new features from a small number of labeled training samples in a fully supervised manner may lead to the over-fitting problem. To solve these problems, much research attention has been paid over the last few years. For example, subspace learning methods, such as Principal Component Analysis (PCA), are applied to reduce feature dimensionality prior to LDA, with the goal of removing null space of the within-class scatter matrix. However, this preprocessing step may lose discriminant information which means the consequent projection in the subspace by LDA may not be the best. A number of methods [7,39,15,19,18,35,41] have been proposed to tackle the SSS problem without losing discriminant information. Though the SSS problem is dealt with, the over-fitting problem persists. Increasing the number of labeled training samples would be an ideal solution. However, data labeling in the real world is usually time-consuming and expensive. Considering the huge amount of data without labels, it is extremely difficult to obtain massive and comparable label information. For this reason, semi-supervised feature analysis methods [14,3,20,6], which make use of both labeled and unlabeled data, have been extensively studied in the past. Unfortunately, most of the existing semi-supervised feature learning algorithms ignore the utilization of both manifold structure and local discriminant information.

In this paper, we propose a new kernel-based feature learning method that can learn new features when the labeled information is limited. Some researchers [25,37,36] have pointed out that exploiting the local structures is more effective and efficient than learning the global structures. Besides, the manifold structure of data is another crucial property to be considered in feature learning. In this work, the contribution can be briefly summarized as the utilization of both the manifold structure and the local discriminant information to deal with the shortage of labeled data points. Compared with those representative semi-supervised feature learning methods, our proposed method not only makes labeled data within the same class closer to each other, but also incorporates the local discriminant information into a joint framework. We find that the local discriminative information of the manifold structure is very important for feature learning, especially when the label information is scarce. This is omitted in most of the previous

works on semi-supervised feature learning. In order to exploit the manifold structure and the local discriminant information, we learn a new graph Laplacian. Specifically, for each data point, we define a local clique in which the data point and its $k-1$ geometric neighbors are included. To achieve this point, kNN is used to exploit the intrinsic manifold structure of data. Moreover, we employ a variant of the Fisher criterion to each clique to evaluate the local discriminant information. The sum of all cliques will be integrated into a joint framework as a global integration. In this way, a new graph Laplacian that holds manifold information with local discriminant information can be learned. This method has two-fold advantages: Firstly, since there are only k data points in each clique (normally quite smaller than the dimensionality of data points), the overall computational burden is greatly relieved. This is because calculating an inverse of a $k \times k$ matrix n times is faster than a direct inverse calculation on a $n \times n$ matrix when n is big. Secondly, it is easy to extend the local discriminant model to a kernel-based version. In this work, we also extend our proposed model using kernel method to learn both labeled and unlabeled data in a high-dimensional space in which data are linearly separable. Additionally, we have proposed an algorithm to solve the optimization problem.

The rest of paper is organized as follows: Notations and definitions will be presented in Section 2. Our proposed method and the optimization algorithm will be elaborated upon in Section 3. In Section 4, we will report the experimental settings, results, and related analysis. The conclusion will be given in the last section.

2 Notations and Definitions

To give a better understanding of the proposed algorithm, notations and definitions used in this paper are summarized in this section. Matrices and vectors are written as boldface uppercase letters and boldface lowercase letters, respectively. In this paper, we follow the conventional definition in semi-supervised learning. In the training dataset, there are n data samples including m labeled data and $n-m$ unlabeled data. Thus, the training data matrix is defined as:

$$\boldsymbol{X} = \left[\boldsymbol{X}_l \; \boldsymbol{X}_u \right] \tag{1}$$

$\boldsymbol{X} \in \mathbb{R}^{d \times n}$. The labeled data are denoted by $\boldsymbol{X}_l \in \mathbb{R}^{d \times m}$, and the unlabeled data are denoted by $\boldsymbol{X}_u \in \mathbb{R}^{d \times (n-m)}$. d is feature dimensionality. Correspondingly, label assignment matrix $\boldsymbol{Y} \in \mathbb{R}^{n \times c}$ is defined as:

$$\boldsymbol{Y} = \begin{bmatrix} \boldsymbol{Y}_l \\ \boldsymbol{Y}_u \end{bmatrix} \tag{2}$$

$\boldsymbol{Y}_l \in \mathbb{R}^{m \times c}$ is for the labeled data while $\boldsymbol{Y}_u \in \mathbb{R}^{(n-m) \times c}$ is for the unlabeled data. c is the number of classes. $\boldsymbol{Y}_u$ is initialized with all zeros and its entry, $y_{it} \in [0, 1]$, denotes how likely the i-th unlabeled data belongs to the t-th class, where $m+1 \leq i \leq n$ and $1 \leq t \leq c$.

3 Locally Discriminative Structure Uncovering

3.1 Proposed method

Inspired by kernel methods, we assume that, after a non-linear mapping function $\boldsymbol{f} = \boldsymbol{\phi}(\boldsymbol{x})$, the mapped data within the same class are still geometrically close to each other in a high-dimensional space that $\boldsymbol{\phi}(\boldsymbol{x})$ maps to. The problem can be formulated as:

$$\min_{\phi} \sum_{t=1}^{c} \sum_{\boldsymbol{x}_p, \boldsymbol{x}_q \in \pi_t} ||\boldsymbol{\phi}(\boldsymbol{x}_p) - \boldsymbol{\phi}(\boldsymbol{x}_q)||^2 \tag{3}$$

π_t contains all the data points from the t-th class. Without loss of generality, we can have:

$$\boldsymbol{\phi}(\boldsymbol{x}) = \sum_{i=1}^{m} \alpha_i \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{x}_i), \tag{4}$$

$\boldsymbol{k}$ is a kernel defined on $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m \in \mathcal{X}$. Note that Eq. (3) needs class information, thus only m labeled data points are considered in Eq. (3). After substituting (4) into (3) by a matrix representation, the objective becomes:

$$\begin{aligned} &\min_{\boldsymbol{\alpha}} \sum_{t=1}^{c} \sum_{\boldsymbol{x}_p, \boldsymbol{x}_q \in \pi_t} ||\boldsymbol{\alpha}^T \boldsymbol{k}(\boldsymbol{X}_l^T, \boldsymbol{x}_p) - \boldsymbol{\alpha}^T \boldsymbol{k}(\boldsymbol{X}_l^T, \boldsymbol{x}_q)||^2 \\ &= \min_{\boldsymbol{\alpha}} \boldsymbol{\alpha}^T \boldsymbol{K}_l \boldsymbol{L}_w \boldsymbol{K}_l \boldsymbol{\alpha}, \end{aligned} \tag{5}$$

where

$$\boldsymbol{L}_w = \boldsymbol{D} - \boldsymbol{W} \tag{6}$$

$\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_m]^T \in \mathbb{R}^m$ is a vector. The weight matrix $\boldsymbol{W}_l \in \mathbb{R}^{m \times m}$ for labeled data is defined as:

$$\boldsymbol{W}_{ij} = \begin{cases} 1 & x_i \text{ and } x_j \text{ are in the same class;} \\ 0 & \text{otherwise,} \end{cases} \tag{7}$$

$\boldsymbol{D}$ is a diagonal matrix with $\boldsymbol{D}_{ii} = \sum_{j=1}^{m} \boldsymbol{W}_{ij}$. $\boldsymbol{K}_l \in \mathbb{R}^{m \times m}$ is the kernel matrix of labeled samples. Note that the representation of $\boldsymbol{\alpha}$ in (5) is an m-dimensional vector that transforms the original feature into one dimension. Now, we extend the learned feature into r dimensions by using a transformation matrix $\boldsymbol{a} = [\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_r] \in \mathbb{R}^{m \times r}$. Consequently, the objective function in (5) is equivalent to

$$\min_{\boldsymbol{a}^T \boldsymbol{K}_l \boldsymbol{a} = I} Tr(\boldsymbol{a}^T \boldsymbol{K}_l \boldsymbol{L}_w \boldsymbol{K}_l \boldsymbol{a}), \tag{8}$$

$Tr(\cdot)$ is trace operator. The constraint $\boldsymbol{a}^T \boldsymbol{K}_l \boldsymbol{a} = I$ is to make the projection orthogonal.

The objective in (8) is based on label information. We aim to make use of both labeled and unlabeled data to improve the performance when the class information is in scarcity. We now extend it into a semi-supervised method by extending the weighted matrix $\boldsymbol{W}_l$ in (7) into a new one:

$$\boldsymbol{W} = \begin{bmatrix} \boldsymbol{W}_l & \boldsymbol{0}_{m \times (n-m)} \\ \boldsymbol{0}_{(n-m) \times m} & \boldsymbol{0}_{(n-m) \times (n-m)} \end{bmatrix} \tag{9}$$

$\boldsymbol{W}_l$ is the weighted matrix for labeled data. $\mathbf{0}_{m\times(n-m)}$ is a zero matrix with m rows and $(n-m)$ columns. In this way, the adjacency graph $\boldsymbol{L}_w$ is enlarged from $m \times m$ to $n \times n$.

As mentioned before, we find that local discriminant information among data, labeled and unlabeled, is quite useful for learning new features especially when the label information is limited. In this work, we aim to learn such local structures and to embed the structures into a joint framework. The objective function in (8) can be re-written as:

$$\min_{\boldsymbol{a}^T\boldsymbol{K}\boldsymbol{a}=I} Tr(\boldsymbol{a}^T\boldsymbol{K}\boldsymbol{L}_w\boldsymbol{K}\boldsymbol{a}) + \lambda\Omega(\cdot), \tag{10}$$

$\Omega(\cdot)$ is the regularization term that exploits the local structures among data samples. λ is the regularization parameter. We assume that all data within the same class are put together and define a scaled class assignment matrix $\boldsymbol{G} \in \mathbb{R}^{n\times c}$ as follows:

$$\boldsymbol{G} = \boldsymbol{Y}(\boldsymbol{Y}^T\boldsymbol{Y})^{-\frac{1}{2}} \tag{11}$$

$\boldsymbol{Y} = [\boldsymbol{y_1}, \ldots, \boldsymbol{y_n}]^T \in \mathbb{R}^{n\times c}$. To exploit the local structure, we define a clique of $\boldsymbol{x}_i$, denoted as $\boldsymbol{\mathcal{N}}_k(\boldsymbol{x}_i)$, which has k data samples containing $\boldsymbol{x}_i$ itself and its $k-1$ neighbors. Given the data matrix $\boldsymbol{X} \in \mathbb{R}^{d\times n}$, the local data matrix for the i-th data is defined as $\boldsymbol{X}_i = [\boldsymbol{x}_i, \boldsymbol{x}_{i_1}, \ldots, \boldsymbol{x}_{i_{k-1}}] \in \mathbb{R}^{d\times k}$. Correspondingly, the local scaled classification matrix for $\boldsymbol{x}_i$ can be defined as $\boldsymbol{G}_i = [\boldsymbol{g}_i, \ldots, \boldsymbol{g}_{i_{k-1}}]^T \in \mathbb{R}^{k\times c}$. Note that both $\boldsymbol{X}_i$ and $\boldsymbol{G}_i$ are actually selected from $\boldsymbol{X}$ and $\boldsymbol{G}$ respectively. We define a selection matrix $\boldsymbol{S}_i \in \mathbb{R}^{n\times k}$ for $\boldsymbol{x}_i$ where $\boldsymbol{S}_i^{pq} = 1$, if $\boldsymbol{x}_p$ is the q-th element of $\boldsymbol{\mathcal{N}}_k(\boldsymbol{x}_i)$, $\boldsymbol{S}_i^{pq} = 0$ otherwise. $1 \le q \le k$. Thus, the local scaled classification matrix for each data can be re-written as $\boldsymbol{G}_i = \boldsymbol{S}_i^T\boldsymbol{G} = \boldsymbol{S}_i^T\boldsymbol{Y}(\boldsymbol{Y}^T\boldsymbol{Y})^{-\frac{1}{2}}$.

In linear discriminant analysis, the objective is to maximize the separability of all data points and to minimize the distances among data that are from the same class. According to the definitions of scatter matrix in linear discriminant analysis, the corresponding total scatter and between class scatter matrices for the local data clique can be defined as $\boldsymbol{S}_{t_i}$ and $\boldsymbol{S}_{b_i}$. To simplify the formulation, we centralize each local data matrix $\bar{\boldsymbol{X}}_i = \boldsymbol{X}_i\boldsymbol{H}$, where $\boldsymbol{H} = \boldsymbol{I}_{k\times k} - \frac{1}{k}\mathbf{1}_k\mathbf{1}_k^T$. $\mathbf{1}_k$ is a k-dimensional column vector with all ones. Thus, we have:

$$\begin{aligned}
\boldsymbol{S}_{t_i} &= \sum_{i=1}^{n}(\boldsymbol{x}_i - \boldsymbol{\mu})(\boldsymbol{x}_i - \boldsymbol{\mu})^T = \bar{\boldsymbol{X}}_i\bar{\boldsymbol{X}}_i{}^T \\
\boldsymbol{S}_{b_i} &= \sum_{t=1}^{c} n_t(\boldsymbol{\mu}_t - \boldsymbol{\mu})(\boldsymbol{\mu}_t - \boldsymbol{\mu})^T = \bar{\boldsymbol{X}}_i\boldsymbol{G}_i\boldsymbol{G}_i^T\bar{\boldsymbol{X}}_i{}^T
\end{aligned} \tag{12}$$

where $\boldsymbol{\mu}_t$ is the mean of samples in the t-th class and $\boldsymbol{\mu}$ is the global mean that is zero after the centralization. n_t is the number of data points of the t-th class. Inspired by the Fisher criterion [8], the optimal local scaled class assignment matrix $\boldsymbol{G}_i^*$ can be obtained by optimizing the follow objective:

$$\begin{aligned}
\boldsymbol{G}_i^* &= \arg\max_{\boldsymbol{G}_i} Tr\left((\boldsymbol{S}_{t_i} + \theta\boldsymbol{I})^{-1}\boldsymbol{S}_{b_i}\right) \\
&= \arg\max_{\boldsymbol{G}_i} Tr\left(\boldsymbol{G}_i^T\bar{\boldsymbol{X}}_i{}^T(\bar{\boldsymbol{X}}_i\bar{\boldsymbol{X}}_i{}^T + \theta\boldsymbol{I})^{-1}\bar{\boldsymbol{X}}_i\boldsymbol{G}_i\right),
\end{aligned} \tag{13}$$

where $\theta \boldsymbol{I}(\theta > 0)$ is added to avoid $(\boldsymbol{S}_{t_i} + \theta \boldsymbol{I})$ be singular. In Eq. (13), $\boldsymbol{G}_i^*$ is a score that evaluates the local discriminant information of each data points. A larger value indicates that the samples in the local clique from different classes are better separated. To control the capacity of local discriminant model, we add a regularization term $Tr(\boldsymbol{G}_i^T \boldsymbol{H} \boldsymbol{G}_i)$. Then the optimization problem in (13) is equivalent to

$$\arg\min_{\boldsymbol{G}_i} Tr\left(\boldsymbol{G}_i^T \boldsymbol{H} \boldsymbol{G}_i - \boldsymbol{G}_i^T \bar{\boldsymbol{X}}_i{}^T (\bar{\boldsymbol{X}}_i \bar{\boldsymbol{X}}_i{}^T + \theta \boldsymbol{I})^{-1} \bar{\boldsymbol{X}}_i \boldsymbol{G}_i\right) \tag{14}$$

It is proved in the supplemental document, the problem of (14) is equivalent to

$$\begin{aligned} &\arg\min_{\boldsymbol{G}_i} Tr\left(\boldsymbol{G}_i^T \boldsymbol{L}_i \boldsymbol{G}_i\right), \\ &\text{where } \boldsymbol{L}_i = \boldsymbol{H}(\bar{\boldsymbol{X}}_i{}^T \bar{\boldsymbol{X}}_i + \theta \boldsymbol{I})^{-1} \boldsymbol{H} \end{aligned} \tag{15}$$

Because of $\boldsymbol{G}_i = \boldsymbol{S}_i^T \boldsymbol{G}$, we take all local manifold structures into account together by summing (15) over all local cliques. Then, the global local discriminant score can be written as:

$$\arg\min_{\boldsymbol{G}} \sum_{i=1}^{n} Tr(\boldsymbol{G}_i^T \boldsymbol{L}_i \boldsymbol{G}_i) = \arg\min_{\boldsymbol{G}} Tr(\boldsymbol{G}^T \boldsymbol{L} \boldsymbol{G}), \tag{16}$$

where

$$\boldsymbol{L} = \sum_{i=1}^{n} \boldsymbol{S}_i \boldsymbol{L}_i \boldsymbol{S}_i^T \tag{17}$$

By using the graph Laplacian in (17), local discriminant manifold structures among data points can be therefore embedded into a joint framework. The objective function in (10) arrives at:

$$\begin{aligned} &\min_{\boldsymbol{a}^T \boldsymbol{K} \boldsymbol{a} = \boldsymbol{I}} Tr(\boldsymbol{a}^T \boldsymbol{K} \boldsymbol{L}_w \boldsymbol{K} \boldsymbol{a}) + \lambda Tr(\boldsymbol{a}^T \boldsymbol{K} \boldsymbol{L} \boldsymbol{K} \boldsymbol{a}) \\ =&\min_{\boldsymbol{a}^T \boldsymbol{K} \boldsymbol{a} = I} Tr(\boldsymbol{a}^T \boldsymbol{K} (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{K} \boldsymbol{a}) \end{aligned} \tag{18}$$

where λ is a parameter that leverages the proportion of utilization of both manifold structure and local discriminant information in the joint framework.

3.2 Optimization

To solve the objective function in (18), we assume

$$\boldsymbol{K} = \boldsymbol{V} \boldsymbol{\Lambda} \boldsymbol{V}^T, \tag{19}$$

and $\tilde{\boldsymbol{V}}$ is the null space of $\boldsymbol{V}$. For any $\boldsymbol{a}$, we can represent

$$\boldsymbol{a} = \boldsymbol{V} \boldsymbol{\beta} + \tilde{\boldsymbol{V}} \boldsymbol{\gamma}, \tag{20}$$

Thus,

$$\begin{aligned} \boldsymbol{a}^T \boldsymbol{K} \boldsymbol{a} &= (\boldsymbol{V} \boldsymbol{\beta} + \tilde{\boldsymbol{V}} \boldsymbol{\gamma})^T \boldsymbol{K} (\boldsymbol{V} \boldsymbol{\beta} + \tilde{\boldsymbol{V}} \boldsymbol{\gamma}) \\ &= \boldsymbol{\beta}^T \boldsymbol{V}^T \boldsymbol{K} \boldsymbol{V} \boldsymbol{\beta} = \boldsymbol{\beta}^T \boldsymbol{\Lambda} \boldsymbol{\beta}, \end{aligned} \tag{21}$$

Algorithm 1: Local Discriminant Structure Uncovering.

Input:

The training data matrix $\boldsymbol{X} = [\boldsymbol{X}_l, \boldsymbol{X}_u] \in \mathbb{R}^{d \times n}$

The classification assignment matrix $\boldsymbol{Y} = [\boldsymbol{Y}_l, \boldsymbol{Y}_u] \in \mathbb{R}^{n \times c}$

Output:

Feature transformation matrix $\boldsymbol{a} \in \mathbb{R}^{n \times r}$

1: Compute kernel matrix $\boldsymbol{K}$;
2: Compute graph $\boldsymbol{L}_w$ and $\boldsymbol{L}$ according to (6) and (17);
3: Compute $\boldsymbol{V}$ and $\boldsymbol{\Lambda}$ by performing eigen-decomposition on $\boldsymbol{K}$ according to (19);
4: Compute $\boldsymbol{\omega}$ by performing eigen-decomposition according to (24);
5: Compute $\boldsymbol{\beta}$ according to (23);
6: Obtain $\boldsymbol{a}$ according to (20);

Substituting (20) and (21) into (18), the objective function is re-formulated as

$$\begin{aligned} &\min_{\boldsymbol{\beta}^T \boldsymbol{\Lambda} \boldsymbol{\beta} = \boldsymbol{I}} Tr(\boldsymbol{a}^T \boldsymbol{K} (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{K} \boldsymbol{a}) \\ = &\min_{\boldsymbol{\beta}^T \boldsymbol{\Lambda} \boldsymbol{\beta} = \boldsymbol{I}} Tr(\boldsymbol{\beta}^T \boldsymbol{\Lambda} \boldsymbol{V}^T (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{V} \boldsymbol{\Lambda} \boldsymbol{\beta}), \end{aligned} \tag{22}$$

Note that $\boldsymbol{\Lambda}$ is invertible, so the solution $\boldsymbol{\beta}$ is the eigenvectors corresponding to $\boldsymbol{\Lambda}^T \boldsymbol{V}^T (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{V} \boldsymbol{\Lambda}$. To enable the solution in the real domain, we can make

$$\boldsymbol{\beta} = \boldsymbol{\Lambda}^{-\frac{1}{2}} \boldsymbol{\omega} \tag{23}$$

and the problem becomes

$$\min_{\boldsymbol{\omega}^T \boldsymbol{\omega} = \boldsymbol{I}} Tr(\boldsymbol{\omega}^T \boldsymbol{\Lambda}^{\frac{1}{2}} \boldsymbol{V}^T (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{V} \boldsymbol{\Lambda}^{\frac{1}{2}} \boldsymbol{\omega}) \tag{24}$$

The optimal solution $\boldsymbol{\omega}$ is the eigenvectors of $\boldsymbol{\Lambda}^{\frac{1}{2}} \boldsymbol{V}^T (\boldsymbol{L}_w + \lambda \boldsymbol{L}) \boldsymbol{V} \boldsymbol{\Lambda}^{\frac{1}{2}}$ with respect to its eigenvalues in an ascending order. We summarize the entire procedure in Algorithm 1 to learn the new features. For any test data $\boldsymbol{x}'$, its l-th new feature is obtained by $\sum_{i=1}^{n} \alpha_{il} \boldsymbol{k}(\boldsymbol{x}', \boldsymbol{x}_i)$, $1 \leq l \leq r$.

4 Experiments

In this section, we will briefly introduce the datasets and the compared methods which are used in the experiments. Afterwards, experimental results are evaluated and analyzed.

4.1 Datasets and Compared Methods

To evaluate our algorithm, we have conducted extensive experiments and compared with a number of approaches on five datasets:

– **COIL-20 [22]**: It contains 1,440 gray-scale images of 20 objects (72 images per object) under various poses. The objects are rotated through 360 degrees and taken at the interval of 5 degrees.
– **UMIST [12]**: The UMIST, which is also known as the Sheffield Face Database, consists of 564 images of 20 individuals. Each individual is shown in a variety of poses from profile to frontal views.
– **USPS [16]**: This dataset collects 9,298 images of handwritten digits (0-9) from envelops by the U.S. Postal Service. All images have been normalized to the same size of 16×16 pixels in grayscale.
– **Yale [9]**: It consists of 2,414 frontal face images of 38 subjects. Different lighting conditions have been considered in this dataset. All images are reshaped into 24×24 pixels.
– **MIMIC II [26]**: It consists of 32,536 patient records in Intensive Care Unit (ICU) at the Beth Israel Deaconess Medical Center (BIDMC) collected from 2001 to 2008. We only extract medical notes from the database together with mortality information (if the patient has been expired at ICU for the first admission). A similar feature extraction pipeline used in [10] has been applied. Differently, Bag-of-Words model is used to encode multiple notes at different times for each patient. Empirically, we set the size of the dictionary as 500. Afterwards, we randomly select 1,000 adult patients of which are positive and negative evenly.

For the first four datasets (COIL20, UMIST, USPS and YaleB), we just simply use their pixels as the input features and the typical RBF kernel. For MIMIC II dataset, we use a 500 dimensional Bag-of-Words representation and a χ^2 kernel. There are a number of kernel functions that have been invented so far. In this paper, our focus is to demonstration the effectiveness of our proposed algorithm rather than making comparisons between different kernels regarding classification performance.

To evaluate our proposed algorithm, we choose several methods as the baseline:

– **Kernel Discriminant Analysis (KDA)** [27]: As one of the representative feature dimensionality reduction methods, KDA aims to project data into a direction on which class centers are far from each other while data samples of the same class are close to each other after feature mapping. We use a speed-up version implemented in [4].
– **Kernel Principal Component Analysis (KPCA)** [28]: Compared to KDA, KPCA reduces feature dimensionality by transforming data into a new coordinate system where the top n greatest variances of data correspondingly lie on the first n coordinates in the new subspace. We test different dimensionality reduction scenarios by KPCA across all the datasets. The best classification performance results are reported.
– **Kernel Semi-supervised Discriminant Analysis (KSDA)** [3]: SDA aims to solve the problem of scarcity of label information when performing discriminant analysis. To utilize unlabeled samples, a graph Laplacian is built to approximate the local geometry of the data manifold where both the labeled and unlabeled data reside. Kernel SDA (KSDA) is used in the experiments.
– **Kernel Semi-supervised Local Fisher discriminant analysis (KSELF)** [30]: SELF which leverages supervised Local Fisher Discriminant Analysis and unsupervised

Principal Component Analysis, is a linear semi-supervised dimensionality reduction method which makes feature analysis effective when only a small number of labeled samples are available. In the experiments, we use its non-linear extension termed as KSELF.

– **Kernel Locality Preserving Projections (KLPP)** [14]: KLPP is an unsupervised manifold learning method which preserves the local structure of samples, i.e. neighborhood relationship, in the original feature space as well as in the new projected space.
– **Co-regularized Ensemble for Feature Selection (EnFS)** [13]: This method employs a co-regularized framework in which a joint $\ell_{2,1}$-norm of multiple feature selection matrices can alleviate the over-fitting problem when the number of labeled data is small. Furthermore, a subset of feature that is more distinctive can be uncovered by removing irrelevant or noisy features.

For all the methods, corresponding parameters are tuned in the same range of $\{10^{-4}$, 10^{-3}, 10^{-2}, 1, 10^{2}, 10^{3}, $10^{4}\}$. Support Vector Machine (SVM) with a linear kernel has been applied as a classifier evaluating the classification performance of each method. The SVM parameter that controls the trade-off between the margin and the size of the slack variables is also tuned in the aforementioned range. The detailed dataset partition is followed by the convention of the semi-supervised learning approaches. Specifically, the training set contains both labeled and unlabeled data, and the testing set is not available during the training phrase. c is denoted as the number of classes for each dataset. In the training dataset, we randomly sample a number of labeled data per class (1, 3, 5, and 10) as different class settings. Therefore the numbers of labeled training data are $1 \times c$, $3 \times c$, $5 \times c$, and $10 \times c$ in the different class settings, while the remaining training data are treated as unlabeled. Particularly, we also investigate conditions in which more labeled information is available by using $25 \times c$, $50 \times c$, $75 \times c$, and $100 \times c$ on MIMIC II dataset. We repeat the experiments five times using the data partitions mentioned above, and report the average results. Because we focus on classification performance, we use Mean Average Precision (MAP) as the effectiveness measure in the comparisons.

4.2 Evaluation

We have made comparisons among all the methods mentioned above over five datasets under different label settings. Generally speaking, our algorithm consistently achieves better classification performance under different settings than all the counterparts. Specifically, from Tab. 1, we can observe that our approach outperforms all supervised and unsupervised methods, including KDA, KPCA and EnFS, when labeled data are quite few, i.e. $1 \times c$. For instance, a number of relatively large margins can be observed on UMIST and USPS. Compared to the semi-supervised methods (KSDA, KSELF) and the unsupervised manifold learning method (KLPP), our method still has superior performance over all the datasets under different settings. In the conditions where more labels are available, it is observed that our method still performs better than all the compared approaches. However, the difference margins between our method and the other approaches are quite limited. For example, our method, on the MIMIC II dataset, performs quite similarly to EnFS and KSELF.

Table 1. Performance comparison (Mean Average Precision ± STD) across all datasets with a linear SVM classifier.

Dataset	Settings	KDA	KPCA	KSDA	KSELF	KLPP	EnFS	Ours
COIL-20	$1 \times c$	0.728±0.026	0.699±0.026	0.777±0.016	0.702±0.032	0.780±0.020	0.702±0.012	**0.818±0.001**
	$3 \times c$	0.832±0.013	0.833±0.018	0.840±0.020	0.819±0.014	0.822±0.007	0.825±0.025	**0.850±0.005**
	$5 \times c$	0.888±0.015	0.881±0.020	0.878±0.022	0.877±0.014	0.876±0.015	0.880±0.013	**0.897±0.002**
	$10 \times c$	0.948±0.009	0.949±0.012	0.926±0.015	0.931±0.012	0.932±0.007	0.935±0.012	**0.957±0.001**
UMIST	$1 \times c$	0.574±0.019	0.558±0.021	0.642±0.026	0.573±0.018	0.580±0.020	0.558±0.017	**0.654±0.010**
	$3 \times c$	0.882±0.033	0.851±0.035	0.860±0.034	0.832±0.044	0.812±0.043	0.810±0.037	**0.889±0.001**
	$5 \times c$	0.960±0.013	0.942±0.019	0.945±0.020	0.937±0.019	0.913±0.021	0.909±0.031	**0.963±0.001**
	$10 \times c$	0.995±0.004	0.989±0.003	0.987±0.008	0.990±0.004	0.967±0.006	0.9779±0.01	**0.996±0.002**
USPS	$1 \times c$	0.536±0.079	0.464±0.056	0.653±0.088	0.351±0.023	0.789±0.056	0.523±0.053	**0.801±0.009**
	$3 \times c$	0.727±0.017	0.720±0.019	0.811±0.015	0.625±0.073	0.905±0.012	0.715±0.027	**0.936±0.008**
	$5 \times c$	0.795±0.024	0.788±0.032	0.865±0.031	0.732±0.028	0.939±0.017	0.788±0.024	**0.948±0.007**
	$10 \times c$	0.868±0.018	0.860±0.007	0.911±0.007	0.837±0.011	0.957±0.005	0.864±0.016	**0.964±0.003**
YaleB	$1 \times c$	0.359±0.013	0.358±0.015	0.245±0.007	0.220±0.011	0.354±0.028	0.216±0.014	**0.463±0.004**
	$3 \times c$	0.872±0.023	0.761±0.033	0.555±0.034	0.652±0.033	0.572±0.033	0.567±0.083	**0.891±0.001**
	$5 \times c$	0.951±0.003	0.907±0.006	0.769±0.021	0.856±0.008	0.717±0.020	0.723±0.039	**0.964±0.002**
	$10 \times c$	0.981±0.003	0.978±0.005	0.948±0.005	0.963±0.011	0.826±0.018	0.863±0.032	**0.992±0.007**
MIMIC II	$25 \times c$	0.683±0.029	0.675±0.028	0.666±0.051	0.699±0.017	0.683±0.037	0.696±0.028	**0.700±0.027**
	$50 \times c$	0.727±0.017	0.722±0.026	0.705±0.035	0.733±0.014	0.712±0.004	0.740±0.008	**0.741±0.018**
	$75 \times c$	0.741±0.012	0.749±0.027	0.727±0.024	0.751±0.010	0.737±0.020	0.754±0.028	**0.755±0.015**
	$100 \times c$	0.754±0.025	0.757±0.021	0.759±0.019	0.757±0.028	0.751±0.020	0.762±0.023	**0.766±0.017**

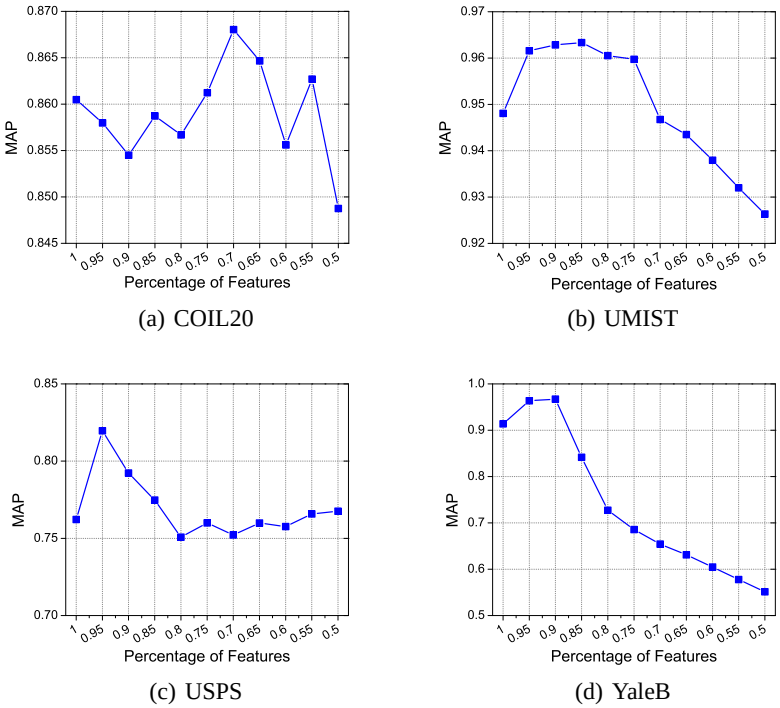

(a) COIL20 (b) UMIST

(c) USPS (d) YaleB

Fig. 1. Performance variations w.r.t dimensionality reduction parameter r.

Apart from the overall classification performance comparisons in Tab. 1, we also studied the effects of parameters used in our method. Due to the page limit, we only studied the sensitivity of parameters on four image datasets. Note that some of the parameters are fixed for demonstration in the following experiments. Thus, the performance results are not as good as the ones in Tab. 1 in which all parameters are tuned to achieve the best performance. In the first round, we test parameter sensitivity for θ and k which are both required when constructing the new graph Laplacian. We test θ in the range of $\{10^{-4}, 10^{-3}, 10^{-2}, 1, 10^{2}, 10^{3}, 10^{4}\}$ and k in the range of $[1, 3, 5, 10, 15, 20]$. We find that the system is not sensitive to θ and the optimal k should be 3 or 5.

In the second round, we firstly fix the aforementioned two parameters of graph Laplacian, $\theta = 1$ and $k = 3$, respectively. Moreover, the regularization parameter, λ, in (24) is set to 1. We plot classification performance changes over four datasets when dimensionality reduction parameter, r, varies. Our aim is to understand how the new features with reduced dimensionality impact performance by fixing all the parameters except the dimension of the inferred feature. The results show that we do not have to use full dimensional features. For example, in Fig. 1(a), the performance scores 86.04% when 100% learned features are preserved on COIL20. It then peaks at 86.67% when 70% are preserved. For each dataset, a further improvement has been observed after

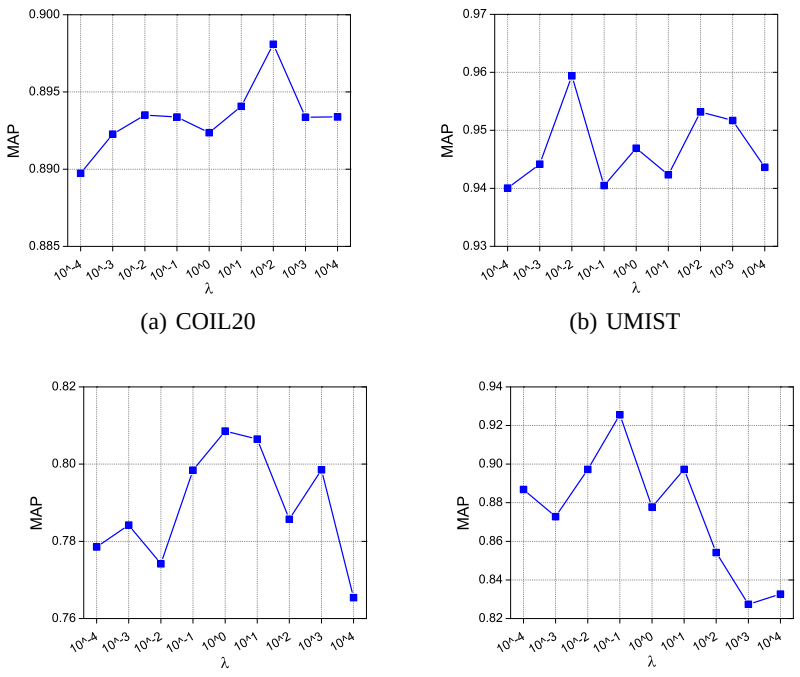

(a) COIL20 (b) UMIST

(c) USPS (d) YaleB

Fig. 2. Performance variations w.r.t regularization parameter λ.

reducing dimensionality of the new features. This improvement might be because irrelevant features and noises are removed by our method after feature mapping in the higher dimensional space.

In the last experiment, all three parameters are fixed ($\theta = 1$, $k = 3$ and $r = 1$). We compare the variations of classification performance when changing the regularization parameter, λ, which leverages local manifold structures in the framework. Note that there is no contribution from local discriminant structure analysis when λ is close to zero. From Fig. 2, the performance on each dataset is a relatively lower value when little local manifold information has been considered. With the variations of λ, for each dataset, the performance varies and scores the best when the weight of the graph Laplacian is increased to a certain amount which is obviously greater than 0. For example, in Fig. 2(b), the performance starts around 94% and almost peaks at 96% when $\lambda = 0.01$, with a nearly 2% improvement. This result confirms that our algorithm successfully incorporates local manifold information into the feature analysis procedure.

5 Conclusion

In this paper, we have proposed a semi-supervised feature analysis method. Specifically, our method enforces data from the same class to become closer to each other in a high-dimensional space after feature mapping. In order to take both local discriminant information and manifold structure into account, a local discriminant model has been applied to the local clique of each data point. Our method successfully learns both labeled and unlabeled data via leveraging the new graph Laplacian that holds local discriminant information. It has proven that our method effectively learns features when the number of labeled data points is quite small.

Acknowledgement

This work was supported by Australian Research Council Discovery Project. The project number is DP140100104. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Australian Research Council.

References

1. Bekios-Calfa, J., Buenaposada, J.M., Baumela, L.: Revisiting linear discriminant techniques in gender recognition. IEEE Trans. Pattern Anal. Mach. Intell. 33(4), 858–864 (2011)
2. Belhumeur, P.N., Hespanha, J.P., Kriegman, D.: Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. IEEE Trans. Pattern Anal. Mach. Intell. 19(7), 711–720 (1997)
3. Cai, D., He, X., Han, J.: Semi-supervised discriminant analysis. In: ICCV (2007)
4. Cai, D., He, X., Han, J.: Speed up kernel discriminant analysis. VLDBJ 20(1), 21–33 (2011)
5. Chang, X., Nie, F., Wang, S., Yang, Y., Zhou, X., Zhang, C.: Compound rank-k projections for bilinear analysis. IEEE Trans. Neural Netw. Learning Syst. 27(7), 1502–1513 (2016)

6. Chang, X., Nie, F., Yang, Y., Huang, H.: A convex formulation for semi-supervised multi-label feature selection. In: AAAI (2014)
7. Chen, L.F., Liao, H.Y.M., Ko, M.T., Lin, J.C., Yu, G.J.: A new lda-based face recognition system which can solve the small sample size problem. Pattern Recogn. 33(10), 1713–1726 (2000)
8. Fukunaga, K.: Introduction to statistical pattern recognition. Academic press (1990)
9. Georghiades, A.S., Belhumeur, P.N., Kriegman, D.: From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE Trans. Pattern Anal. Mach. Intell. 23(6), 643–660 (2001)
10. Ghassemi, M., Naumann, T., Doshi-Velez, F., Brimmer, N., Joshi, R., Rumshisky, A., Szolovits, P.: Unfolding physiological state: Mortality modelling in intensive care units. In: SIGKDD. pp. 75–84 (2014)
11. Goudail, F., Lange, E., Iwamoto, T., Kyuma, K., Otsu, N.: Face recognition system using local autocorrelations and multiscale integration. IEEE Trans. Pattern Anal. Mach. Intell. 18(10), 1024–1028 (1996)
12. Graham, D., Allinson, N.: Characterizing virtual eigensignatures for general purpose face recognition. in face recognition: From theory to applications. NATO ASI Series F, Computer and Systems Sciences 163, 446–456 (1998)
13. Han, Y., Yang, Y., Zhou, X.: Co-regularized ensemble for feature selection. In: AAAI. pp. 1380–1386 (2013)
14. He, X., Niyogi, P.: Locality preserving projections. In: NIPS. p. 153 (2004)
15. Huang, R., Liu, Q., Lu, H., Ma, S.: Solving the small sample size problem of lda. In: ICPR. vol. 3, pp. 29–32 (2002)
16. Hull, J.J.: A database for handwritten text recognition research. IEEE Trans. Pattern Anal. Mach. Intell. 16(5), 550–554 (may 1994)
17. Iosifidis, A., Tefas, A., Nikolaidis, N., Pitas, I.: Multi-view human movement recognition based on fuzzy distances and linear discriminant analysis. CVIU 116(3), 347–360 (2012)
18. Jiang, X., Mandal, B., Kot, A.: Eigenfeature regularization and extraction in face recognition. IEEE Trans. Pattern Anal. Mach. Intell. 30(3), 383–394 (2008)
19. Kyperountas, M., Tefas, A., Pitas, I.: Weighted piecewise lda for solving the small sample size problem in face verification. TNN 18(2), 506–519 (March 2007)
20. Li, S., Fu, Y.: Low-rank coding with b-matching constraint for semi-supervised classification. In: AAAI. pp. 1472–1478 (2013)
21. Lu, Z., Peng, Y.: Learning descriptive visual representation by semantic regularized matrix factorization. In: AAAI. pp. 1523–1529 (2013)
22. Murase, H., Nayar, S.K.: Visual learning and recognition of 3-d objects from appearance. IJCV 14(1), 5–24 (1995)
23. Nie, F., Xiang, S., Liu, Y., Hou, C., Zhang, C.: Orthogonal vs. uncorrelated least squares discriminant analysis for feature extraction. Pattern Recognition Letters 33(5), 485–491 (2012)
24. Nie, F., Zeng, Z., Tsang, I.W., Xu, D., Zhang, C.: Spectral embedded clustering: A framework for in-sample and out-of-sample spectral clustering. IEEE Trans. Neural Networks 22(11), 1796–1808 (2011)
25. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. Science 290(5500), 2323–2326 (2000)
26. Saeed, M., Villarroel, M., Reisner, A., Clifford, G., Lehman, L., Moody, G., Heldt, T., Kyaw, T., Moody, B., Mark, R.: Multiparameter intelligent monitoring in intensive care ii (mimic-ii): a public-access intensive care unit database. Crit. Care Med. 39(5), 952 (2011)
27. Schölkopf, B., Müller, K.R.: Fisher discriminant analysis with kernels. NNSP (1999)
28. Schölkopf, B., Smola, A., Müller, K.R.: Nonlinear component analysis as a kernel eigenvalue problem. Neural Comput. 10(5), 1299–1319 (1998)

29. Shu, X., Gao, Y., Lu, H.: Efficient linear discriminant analysis with locality preserving for face recognition. Pattern Recogn. 45(5), 1892–1898 (2012)
30. Sugiyama, M., Idé, T., Nakajima, S., Sese, J.: Semi-supervised local fisher discriminant analysis for dimensionality reduction. Mach. Learn. 78(1-2), 35–61 (2010)
31. Swets, D.L., Weng, J.J.: Using discriminant eigenfeatures for image retrieval. IEEE Trans. Pattern Anal. Mach. Intell. 18(8), 831–836 (1996)
32. Wang, S., Nie, F., Chang, X., Yao, L., Li, X., Sheng, Q.Z.: Unsupervised feature analysis with class margin optimization. In: ECMLPKDD. pp. 383–398. Springer (2015)
33. Wang, S., Pan, P., Long, G., Chen, W., Li, X., Sheng, Q.Z.: Compact representation for large-scale unconstrained video analysis. World Wide Web 19(2), 231–246 (2016)
34. Yang, J., Yu, K., Gong, Y., Huang, T.: Linear spatial pyramid matching using sparse coding for image classification. In: CVPR. pp. 1794–1801 (2009)
35. Yang, Y., Ma, Z., Nie, F., Chang, X., Hauptmann, A.G.: Multi-class active learning by uncertainty sampling with diversity maximization. International Journal of Computer Vision 113(2), 113–127 (2015)
36. Yang, Y., Nie, F., Xu, D., Luo, J., Zhuang, Y., Pan, Y.: A multimedia retrieval framework based on semi-supervised ranking and relevance feedback. IEEE Trans. Pattern Anal. Mach. Intell. 34(4), 723–742 (2012)
37. Yang, Y., Zhuang, Y.T., Wu, F., Pan, Y.H.: Harmonizing hierarchical manifolds for multimedia document semantics understanding and cross-media retrieval. IEEE Trans. Multimedia 10(3), 437–446 (2008)
38. Ye, J., Li, Q.: Lda/qr: an efficient and effective dimension reduction algorithm and its theoretical foundation. Pattern Recogn. 37(4), 851–854 (2004)
39. Yu, H., Yang, J.: A direct lda algorithm for high-dimensional data with application to face recognition. Pattern Recogn. 34, 2067–2070 (2001)
40. Zhang, C., Nie, F., Xiang, S.: A general kernelization framework for learning algorithms based on kernel PCA. Neurocomputing 73(4-6), 959–967 (2010)
41. Zhang, T., Fang, B., Tang, Y.Y., Shang, Z., Xu, B.: Generalized discriminant analysis: A matrix exponential approach. IEEE Trans. Cybern. 40(1), 186–197 (2010)
42. Zhu, X., Li, X., Zhang, S.: Block-row sparse multiview multilabel learning for image classification. IEEE Trans. Cybern. 46(2), 450–461 (2016)
43. Zhu, X., Li, X., Zhang, S., Ju, C., Wu, X.: Robust joint graph sparse coding for unsupervised spectral feature selection. IEEE Trans. Neural Netw. Learn. Syst. (2016)
44. Zhu, X., Suk, H., Lee, S., Shen, D.: Subspace regularized sparse multitask learning for multiclass neurodegenerative disease identification. IEEE Trans. Biomed. Eng. 63(3), 607–618 (2016)
45. Zhu, X., Suk, H.I., Wang, L., Lee, S.W., Shen, D., Alzheimers Disease Neuroimaging Initiative, et al.: A novel relational regularization feature selection method for joint regression and classification in ad diagnosis. Med. Image Anal. (2015)